

Express Mailing Label No.: ER211533156US

PATENT APPLICATION

Docket No.: 1200.2.78

IBM Docket No.: SJO9-2003-0008

UNITED STATES PATENT APPLICATION

of

Richard V. Kisley,

John M. Lake,

and

Durga D. Mannaru

for

**INCREMENTAL DATA STORAGE
METHOD, APPARATUS, INTERFACE, AND SYSTEM**

INCREMENTAL DATA STORAGE METHOD, APPARATUS, INTERFACE, AND SYSTEM

BACKGROUND OF THE INVENTION

1. The Field of the Invention

The invention relates to interfaces, methods, apparatus, and systems for storing data. Specifically, the invention relates to interfaces, methods, apparatus, and systems for managing incremental data storage.

2. The Relevant Art

Data storage devices and systems are subject to equipment failures, software bugs, operational errors, and disasters that prevent immediate or long term access to valuable data.

Additionally, data storage systems often support environments and applications that require storage reliability and availability that is greater than what is inherently provided by individual storage devices. As a result, many techniques and strategies have been applied to data storage devices and systems in order to increase the reliability and availability of data access. Some of those techniques include storage redundancy, data archiving, and remote copy techniques.

Storage redundancy techniques improve reliability by storing redundant images on multiple devices. Data mirroring, in which data is written to two or more devices in parallel, is perhaps the simplest redundancy technique. Although expensive in terms of storage usage, data mirroring facilitates increased performance in that read operations may be distributed to each parallel device to increase storage throughput. However, storage redundancy techniques, although useful, do not address the need for access to previous copies of data files or images within a data storage system.

1 Data archiving techniques were developed in response to the need to access previous
2 copies of data files or images within data storage systems. With data archiving techniques,
3 data files or images are copied to backup devices and maintained for subsequent use. Backup
4 devices typically comprise slower media that must be carefully managed to successfully
5 restore a file or image to a previous state. To reduce the bandwidth associated with archiving
6 data, incremental techniques are often applied in which only the changes to a file or image
7 are stored on the archive media. Incremental techniques while useful, may increase the
8 complexity of restoring data files or images within data storage systems.

9 Remote copy techniques are a particular form of data archiving in which data is
10 copied to a remote site either via physical media or transmission links. The availability of a
11 remote copy facilitates recovery operations in the event of a local disaster such as a fire or
12 flood. However, providing complete and constant data availability requires transmission of
13 each write operation to the remote site resulting in high bandwidth requirements for the
14 transmission link.

15 The aforementioned techniques and strategies are often implemented as host-based
16 solutions requiring considerable planning, administration, and coordination for reliable
17 deployment. Typically, multiple techniques and strategies must be deployed and managed to
18 fit the needs of a particular installation. Such solutions are typically expensive to implement,
19 reduce system performance, and may not provide adequate data availability and reliability.

20 Consequently, a need for an integrated approach to increasing data availability and
21 reliability that is inexpensive to deploy, easy to manage, and transparent to system
22 performance. In particular, what is needed are interfaces, methods, apparatus, and systems
23 for providing and managing incremental data storage in a manner that supports policy
24 management, data redundancy, and compacted remote copy operations.

SUMMARY OF THE INVENTION

The various elements of the present invention have been developed in response to the present state of the art, and in particular, in response to the problems and needs in the art that have not yet been fully solved by currently available data storage means and methods. Accordingly, the present invention provides an improved interface, method, apparatus, and system for storing and managing incremental copies of volume images.

In one aspect of the present invention, an interface for managing incremental data storage is provided and includes a write function configured to append an entry to an incremental log, a read function configured to retrieve a most recent log entry corresponding to a block address, and a snapshot function configured to automatically partition the incremental log into a first and second volume.

The interface may also include a delete volume function configured to release a snapshot volume, a policy assignment function configured to assign a policy to an incremental log, a read entry function configured to retrieve a sequential entry from the incremental log, and a compact volume function configured to compact a snapshot volume. The interface for managing incremental data storage provides a rich set of functions that facilitate automated snapshot management, compacted or non-compacted remote copy operations, policies for storage management, and advanced data recovery. Snapshot images may be automatically associated with a volume identifier such as a logical unit number (LUN) in order to reduce the complexity of managing snapshot data.

In another aspect of the present invention, a method for managing incremental data storage includes appending data to an incremental log, automatically partitioning the incremental log in response to a snapshot operation, and assigning a volume identifier to a newly formed partition as directed by a storage management policy.

The aforementioned method may also include copying selected log entries to a tertiary volume, retrieving a most recent log entry corresponding to a block address, retrieving a most recent log entry corresponding to a specified snapshot volume and block

1 address, and automatically compacting a snapshot partition. In one embodiment, each log
2 entry contains a pointer to a subsequent entry and in-place compaction of a snapshot partition
3 is performed by placing redundant log entries (i.e. overwritten entries) within a free list. In-
4 place compaction substantially eliminates the copy operations normally associated with
5 compaction.

6 The method for managing incremental data storage provides an integrated approach
7 to managing incremental storage that supports a wide variety of policies such as automated
8 volume creation and automated compaction of incremental partitions into a baseline image in
9 response to snapshot operations. The provided functionality greatly reduces the complexity
10 of managing incremental storage and facilitates increased data storage reliability and
11 availability.

12 In another aspect of the present invention, an apparatus for incremental data storage
13 includes a baseline volume containing a baseline image, an incremental log partitioned into
14 one or more snapshot volumes, and a partition module configured to automatically partition
15 the incremental log into an additional snapshot volume in response to a snapshot operation.
16 In one embodiment, the partition module is also configured to assign a volume identifier to a
17 newly formed partition as directed by a storage management policy.

18 In addition to the aforementioned elements, the apparatus for incremental data storage
19 may also include a copy module configured to copy selected entries from the incremental log
20 to a tertiary volume, a read module configured to retrieve a most recent log entry
21 corresponding to a specified snapshot and block address, and a compaction module
22 configured to automatically compact a snapshot partition to the baseline volume. In one
23 embodiment, the compaction module is configured to conduct in-place compaction. The
24 apparatus for incremental data storage provides incremental data storage and supports
25 transparent storage management within a data processing system.

26 The various aspects and elements of the present invention are combined into a system
27 for redundant incremental data storage. In one embodiment, the system for redundant

1 incremental data storage includes a primary storage device configured to store data, a
2 secondary storage device configured to store data within a baseline volume and an
3 incremental log, the incremental log comprising at least one snapshot volume; a controller
4 configured to store and access data on the primary and secondary storage device, and a
5 snapshot management module configured to automatically partition the incremental log into
6 an additional snapshot volume in response to a snapshot operation. In one embodiment, the
7 primary storage device comprises a plurality of redundantly arranged devices and the
8 secondary storage device provides an additional level of data redundancy. The system for
9 redundant incremental data storage may also include a host configured to access data via the
10 controller.

11 The various elements and aspects of the present invention provide highly available
12 and reliable data storage to a data processing system or the like. These and other features and
13 advantages of the present invention will become more fully apparent from the following
14 description and appended claims, or may be learned by the practice of the invention as set
15 forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

In order that the manner in which the advantages of the invention are obtained will be readily understood, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments thereof, which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

Figure 1 is a block diagram illustrating one embodiment of a prior art data storage system;

Figure 2 is a block diagram illustrating one embodiment of a redundant incremental storage system of the present invention;

Figure 3 is a block diagram illustrating one embodiment of an incremental storage appliance of the present invention;

Figure 4 is a flow chart diagram illustrating one embodiment of an incremental storage management method of the present invention;

Figure 5 is a text-based diagram depicting one embodiment of an incremental storage management interface of the present invention; and

Figure 6 is a flow chart diagram illustrating one embodiment of an in-place compaction method of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Many of the functional units described in this specification have been labeled as modules, in order to more particularly emphasize their implementation independence. For example, modules may be implemented in software for execution by various types of processors. An identified module of executable code may, for instance, comprise one or more physical or logical blocks of computer instructions which may, for instance, be organized as an object, procedure, or function. Nevertheless, the executables of an identified module need not be physically located together, but may comprise disparate instructions stored in different locations which, when joined logically together, comprise the module and achieve the stated purpose for the module. For example, a module of executable code could be a single instruction, or many instructions, and may even be distributed over several different code segments, among different programs, and across several memory devices.

Modules may also be implemented in hardware as electronic circuits comprising custom VLSI circuitry, off-the-shelf semiconductors such as logic chips, transistors, or other discrete components. A module may also be implemented in programmable hardware devices such as field programmable gate arrays, programmable array logic, programmable logic devices or the like.

Similarly, operational data may be identified and illustrated herein within modules, and may be embodied in any suitable form and organized within any suitable type of data structure. The operational data may be collected as a single data set, or may be distributed over different locations including over different storage devices, and may exist, at least partially, merely as electronic signals on a system or network.

Referring to Figure 1, a representative prior art storage system 100 includes a host 110, one or more storage controllers 120 with redundant storage devices 130, one or more backup devices 140, and one or more data vaults 150. The various devices and modules depicted in Figure 1 portray the complexity in providing data storage that is reliable and

1 continuously available in light of equipment failures, operator errors, site disasters, and the
2 like.

3 The depicted storage controller 120 includes a data redundancy module 122. The
4 data redundancy module 122 generates and stores redundant images on the redundant storage
5 devices 130. In certain configurations, the redundancy may be a completely redundant (i.e.,
6 mirrored) image storage stored on more than one device. In such a configuration, data access
7 performance may be increased by alternating read operations between the redundant storage
8 devices 130.

9 The depicted host 110 includes a snapshot management module 112 and a remote
10 storage management module 114. The snapshot management module 112 captures volume
11 images from the redundant storage devices 130 and streams those images to the backup
12 device 140. The snapshot management module 112 may copy only those portions of an
13 image which have changed since a previous backup.

14 In certain embodiments, portions of the snapshot management module may be
15 contained within the storage controller 120 in order to increase system performance.
16 Nevertheless, snapshot management and data redundancy are typically separate functions
17 within a prior art storage system in contrast to the integrated nature of the present invention.

18 The depicted remote storage management module 114 transmits data on a data link
19 116 to the data vault 150. Typically, in order to maintain coherency between the image on
20 the local storage volumes and the replicated volumes on the data vault 150, each write
21 operation must be transmitted to the data vault 150 by the remote storage management
22 module 114. However, due to the high bandwidth associated with write operations, large
23 data buffers and high bandwidth transmission links are required in order to keep up with
24 local write operations.

25 In certain embodiments, physical media may be transported to the data vault 150 in
26 lieu of transmission on the data link 116. However, writing data to archival media increases
27

1 the window of vulnerability in the event of failures and requires tedious synchronization and
2 coordination activities when conducting recovery operations.

3 Figure 2 is a block diagram illustrating one embodiment of a redundant incremental
4 storage system 200 of the present invention. As depicted, the redundant incremental storage
5 system includes a host 210, a storage controller 220, a set of redundant storage devices 230,
6 and the data vault 150. The redundant incremental storage system 200 provides an integrated
7 approach to reliable high-availability data storage resulting in an inexpensive, easy to manage
8 solution that is transparent to system performance.

9 As depicted, the storage controller 220 includes a storage management module 222, a
10 data redundancy module 224, and a remote archiving module 226. Although the modules
11 222, 224, and 226 are depicted within a single storage controller 220, one of skill in the art
12 will appreciate that the functionality provided by the modules 222, 224, and 226 may be
13 distributed across a plurality of storage controllers and devices.

14 The storage management module 222 provides an integrated interface and related
15 functionality for various storage management operations such as snapshot operations,
16 compaction operations, archive operations, volume partitioning operations, and the like. The
17 interface may conform to a storage management method and API described in conjunction
18 with Figures 4 and 5.

19 In certain embodiments, the storage management module 222 supports the
20 aforementioned operations via policies 223 such as temporal-based policies, status-based
21 policies, and event-based policies. Through policies, snapshot operations, compaction
22 operations, archive operations, and the like may be conducted at specific times, at specific
23 intervals, in response to certain events, or conditions such as space availability on a volume.

24 The remote archiving module 226 may transmit compacted or non-compacted,
25 incremental or non-incremental images of volumes on a data link 228 to the data vault 150.
26 In certain embodiments, in place of the data vault 150, a completely redundant site (not
27 shown) may be connected to the storage system 200 via the data link 228.

1 The redundant storage devices 230 include one or more incremental storage devices
2 230b and may include one or more non-incremental storage devices 230a. The incremental
3 storage devices 230b comprise baseline volumes 232 and incremental volumes 234. In
4 certain embodiments, the incremental volumes 234 correspond to partitions within an
5 incremental log (not shown) and each partition contains entries for each write operation that
6 occurred over a particular snapshot interval.

7 In one embodiment, the incremental log is automatically partitioned in response to a
8 snapshot operation, thereby reducing the complexity of managing and accessing point-in-
9 time copies of a particular volume. In certain embodiments, the baseline volumes 232 and
10 the incremental volumes 234 function as secondary volumes to the volumes on the non-
11 incremental storage devices and essentially form a RAID-1 redundancy configuration.

12 The non-incremental storage devices 230a and the incremental storage devices 230b
13 may be arranged in various redundant configurations either separately or in combination. For
14 example, the non-incremental storage devices 230a may be arranged in a RAID-5
15 configuration, while the incremental storage devices 230b may be arranged with little or no
16 redundancy. In such an arrangement, the combination of the non-incremental storage devices
17 230a and the incremental storage devices 230b provide an additional level of redundancy
18 beyond the separate RAID configurations of the non-incremental storage devices 230a and
19 the incremental storage devices 230b.

20 Figure 3 is a block diagram illustrating one embodiment of an incremental storage
21 appliance 300 of the present invention. The depicted incremental storage appliance 300
22 illustrates one embodiment of certain elements of the redundant incremental storage system
23 200 in greater detail. In one embodiment, the depicted elements are integrated into a
24 standalone storage subsystem. In another embodiment, the depicted elements correspond to
25 volumes or partitions within the incremental storage devices 230b and modules residing on
26 the storage controller 220, the host 210, or the like.

1 The incremental storage appliance 300 provides incremental storage in a manner that
2 supports policy management, data redundancy, and compacted remote copy operations. The
3 ability to support policy management, data redundancy, and compacted remote copy
4 operations increases the utility and convenience of the storage appliance 300 over
5 conventional storage means.

6 As depicted, the storage appliance 300 includes the storage management module 222,
7 a read/write module 310, a compaction module 320, an archive module 330, a partition
8 module 340, and a storage medium 350. The storage medium 350 may include a plurality of
9 partitions 352 associated with volumes including baseline volumes 360 and incremental
10 volumes 380 including current volumes 390. In the depicted embodiment, the partitions
11 associated with the incremental volumes 380 are partitioned sequentially within the
12 incremental log 370 as a result, for example, of a snapshot operation.

13 The volumes associated with the storage medium 350 may include a baseline volume
14 360 associated with a baseline partition 362, and several incremental volumes 380 including
15 a current volume 390 associated with a set of incremental partitions 382. As depicted, the
16 partitions associated with the baseline volumes 360 and the incremental volumes 380
17 include, respectively, one or more baseline entries 365, and one or more incremental entries
18 385. The depicted baseline entries 365 and the incremental entries 385 correspond to data
19 blocks provided by the read/write module 310 in conjunction with write operations and may
20 contain a block address to facilitate recovery operations and the like.

21 The read/write module 310 supports read and write operations conducted on the
22 storage medium 350. As depicted, write operations directed to the storage medium 350 are
23 appended to the incremental log 370 and placed within a partition corresponding to the
24 current volume 390 in sequential (*i.e.* time) order. An index table (not shown) or other
25 appropriate mechanism, may be maintained within the read/write module 310 or elsewhere,
26 indicating the placement of the most recent entry corresponding to a block address within a
27

1 specified volume. The index may be updated with each write operation in order to reflect
2 updates to a particular block.

3 In one embodiment, a block address and a volume identifier, such as a logical unit
4 number (LUN), are provided when conducting read operations in order to indicate the block
5 that is to be read within the specified volume. The volume identifier may specify a baseline
6 volume or an incremental volume such as the current volume. In one embodiment, the
7 aforementioned index, or the like, is used to retrieve the most recent entry on the storage
8 medium 350 that corresponds to the specified volume and block. Given the incremental
9 structure of the storage medium 350, the most recent entry may reside in a previous partition
10 such as the baseline partition 362 rather than the partition associated with the specified
11 volume.

12 The compaction module 320 provides compaction services within the storage
13 appliance 300. Compaction may be conducted in conjunction with merging an incremental
14 volume into the baseline volume 360, upon request, or in conjunction with remote copy
15 operations, or the like. In one embodiment, compaction involves reusing space containing
16 redundant entries in a manner that eliminates copy or move operations normally associated
17 with compaction. In another embodiment, compaction involves transferring the most recent
18 entry for each block address to the baseline volume, a backup device, a data vault, a
19 redundant site, or the like.

20 The archive module 330 conducts archiving operations to a backup device or system
21 such as a remote data vault, or the like. The archive module 330 may transmit compacted or
22 non-compacted, incremental or non-incremental, time-ordered or block-ordered images, or
23 combinations thereof, to the designated backup device or system. In one embodiment, the
24 compacted images are generated by the compaction module 320.

25 The partition module 340 manages the incremental partitions 382 within the
26 incremental log 370 and may manage the baseline partition 362. In one embodiment, the
27 partition module 340 automatically partitions the incremental log 370 in response to a

1 snapshot operation. The partition module 340 may also automatically assign a volume
2 identifier to a newly formed partition in order to facilitate convenient access to specific
3 snapshots.

4 The functionality of the compaction module 320, the archive module 330, and the
5 partition module 340 may be managed and accessed by the storage management module 222
6 as described in conjunction with Figure 2. Specifically the storage management module 222
7 may provide an API to access the functionality of the aforementioned modules and to set
8 policies 223 governing their operation such as temporal-based policies, status-based policies,
9 and event-based policies.

10 Figure 4 is a flow chart illustrating one embodiment of an incremental storage
11 management method 400 of the present invention. The incremental storage management
12 method 400 may be conducted in conjunction with, or independent of, the incremental
13 storage appliance 300, the storage controller 220 (specifically the storage management
14 module 222), or the like. The incremental storage management method 400 facilitates
15 conducting snapshot management, remote copy, data compaction, policy management, and
16 other operations associated with data storage devices and systems.

17 As depicted, the incremental storage management method 400 comprises a command
18 loop that includes a receive command step 405, a series of tests 410-480, and a
19 corresponding set of routines 415-485. The depicted series of tests include a read block test
20 410, a write block test 420, a create snapshot test 430, a delete snapshot test 440, a set policy
21 test 450, a read entry test 460, a compact volume test 470 and a shutdown test 480. The
22 commands executed by the incremental storage management method 400 may be sent by an
23 application, utility, or the like, executing on a host such as the host 210 depicted in Figure 2.
24 In response to reception of a command at step 405, the depicted tests are conducted to
25 ascertain the corresponding routine to be conducted by the incremental storage management
26 method 400.

1 The corresponding set of routines include a retrieve block routine 415, an append
2 entry routine 425, a partition log routine 435, a delete volume routine 445, a set policy
3 routine 455, a retrieve entry routine 465, a compact volume routine 475, and an initiate
4 shutdown routine 485. Although depicted as a command loop, other execution structures
5 familiar to those of skill in the art, such as an index-driven interrupt table or a code object,
6 may be employed to embody the incremental storage management method 400.

7 The retrieve block routine 415 retrieves a data block corresponding to a specified
8 block address. Optionally, a specific volume may be specified. In one embodiment, the data
9 block retrieved is the most recent entry within either the baseline partition or an incremental
10 partition that is older than the specified volume.

11 The append entry routine 425 appends a log entry such as a data block corresponding
12 to a write operation to the incremental log. In one embodiment, an index table containing the
13 location of the most recent version of each data block is updated to reflect the new entry
14 within the incremental log.

15 The partition log routine 435 partitions the incremental log and may automatically
16 associate a volume identifier with the newly formed partition. The depicted partition log
17 routine 435 may be conducted in response to reception of a 'create snapshot' command, or
18 the like. In certain embodiments, the value of the volume identifier (such as a LUN) that is
19 associated with the newly formed partition is controlled by policies set by a client
20 application, system utility, system administrator, or the like. For example, depending on the
21 selected policy, the volume identifier of the newly formed partition may be a system defined
22 constant, a sequential value, an explicitly specified value such as a parameter associated with
23 the 'create snapshot' command, or the like.

24 The delete volume routine 445 disassociates a volume identifier with a partition. The
25 delete volume routine 445 may also initiate compaction of an incremental partition into the
26 partition corresponding to the baseline volume. In one embodiment, a storage management
27

1 policy may direct the delete volume routine 445 to place the released volume identifier into
2 an available pool.

3 The set policy routine 455 assigns a specified storage management policy to an
4 explicitly or implicitly specified resource. The specified resource is then managed according
5 to the specified policy. In one embodiment, the set policy routine essentially assigns specific
6 values to various system options and control parameters in order to effect the specified
7 policy. Examples of system options directed by policy assignments include options for
8 automatic compaction of volumes, options for assigning volume identifiers, and options for
9 data redundancy.

10 The retrieve entry routine 465 retrieves an entry from a partition or volume in
11 response to a read entry command or the like. In one embodiment, entries are sequentially
12 accessed in time order. The retrieve entry routine is particularly useful when conducting
13 recovery or other snapshot management operations. For example, the retrieve entry routine
14 may be used to assemble an image of a volume corresponding to a specified time by
15 sequentially accessing the incremental log up to the specified time.

16 The compact volume routine 475 compacts a specified volume. In one embodiment,
17 the volumes may be compacted in place using the method illustrated in Figure 6 or directed
18 to a specific target such as the baseline volume. Other compaction strategies and methods
19 known to those of skill in the art may be deployed by the compact volume routine 475.

20 The initiate shutdown routine 485 is used to shutdown the incremental storage
21 management method 400 and associated hardware. Upon completion of the shutdown
22 routine 485, the incremental storage management method 400 ends 490.

23 The collection of routines provided by the incremental storage management method
24 400 facilitates managing incremental storage in a manner that facilitates automated snapshot
25 management, remote copy operations, policy management, and the like.

26 Figure 5 is a text-based diagram depicting one example of an incremental storage
27 interface 500 of the present invention. In one embodiment, the functions depicted in Figure 5

1 correspond directly to the routines of the incremental storage management method 400
2 described in conjunction with Figure 4. In one embodiment, the incremental storage
3 interface 500 is supported by the storage management module 222.

4 The incremental storage interface 500 facilitates invoking corresponding routines (via
5 messaging, code marshaling, or the like) contained within modules of the incremental storage
6 appliance 300, the storage controller 220, or the like. The corresponding routines may be
7 invoked by an operating system, a utility, an application, or the like residing on a host.
8 Additionally, the incremental storage interface 500 may be invoked by modules resident or
9 co-resident on a device or system such as the various modules depicted in Figures 1-4.

10 As depicted, the incremental storage interface 500 includes a set policy function 510,
11 a read block function 520, a write block function 530, a create snapshot function 540, a
12 delete snapshot function 550, a read snapshot entry function 560, and a compact volume
13 function 570. In one embodiment, the functions 510-570 correspond directly to the routines
14 415-475 described in conjunction with Figure 4.

15 Figure 6 is a flow chart diagram illustrating one embodiment of an in-place
16 compaction method 600 of the present invention. The in-place compaction method 600 may
17 be conducted in conjunction with, or independent of the compaction module 320, the
18 compact volume routine 475, or the compact volume function 570. As depicted, the in-place
19 compaction method 600 includes an initialize data step 610, an inspect entry step 620, a
20 redundant entry test 630, a free entry step 640, and an additional entry test 650. The in-
21 place compaction method 600 facilitates compaction of a snapshot partition, or the like,
22 without moving entries within an incremental log.

23 The initialize data step 610 initializes data related to conducting in-place compaction.
24 In one embodiment, the initialize data step 610 initializes a free list used to record the entries
25 that may be reused, and a block encountered table used to track which blocks have been
26 encountered within a snapshot partition.
27

1 The inspect entry step 620, inspects an entry within a snapshot partition. In certain
2 embodiments, the incremental log entries are inspected in reverse order (of occurrence)
3 beginning with the most recent entry. In one embodiment, inspecting includes retrieving a
4 logical block identifier associated with the entry.

5 The redundant entry test 630 ascertains whether an incremental log entry is redundant
6 and therefore available for reuse. In one embodiment, the redundant entry test 630 includes a
7 test and set operation that accesses a bit corresponding to the logical block identifier within a
8 block encountered table (not shown.) In the aforementioned embodiment, the test operation
9 ascertains whether the block was previously encountered and the set operation indicates (for
10 subsequent tests) that the block has been encountered.

11 If the redundant entry test 630 ascertains that the block has not been encountered, the
12 method loops to the inspect entry step 620. If the block has been encountered (*i.e.* is
13 overwritten and therefore redundant), the method proceeds to the free entry step 640. The
14 free entry step 640 frees an incremental log entry for subsequent reuse. In one embodiment,
15 the free entry step 640 adds the entry to a free list.

16 Subsequent to completion of the free entry step 640, the in-place compaction method
17 600 proceeds to the additional entry test 650. The additional entry test 650 ascertains
18 whether additional entries remain within the partitioned being processed. If additional
19 entries remain, the method loops to the inspect entry step 620. If no additional entries
20 remain, the method ends 660.

21 The present invention eases and improves storage management, particularly
22 management of incremental images within a storage system. The present invention may be
23 embodied in other specific forms without departing from its spirit or essential characteristics.

24 The described embodiments are to be considered in all respects only as illustrative and not
25 restrictive. The scope of the invention is, therefore, indicated by the appended claims rather
26 than by the foregoing description. All changes which come within the meaning and range of
27 equivalency of the claims are to be embraced within their scope.